5 Using Intermediate Design Artifacts to Bridge the Gap Between Cognitive Analysis and Cognitive Engineering

SCOTT S. POTTER
Aegis Research Corporation

WILLIAM C. ELM
Aegis Research Corporation

EMILIE M. ROTH
Roth Cognitive Engineering

JAMES W. GUALTIERI
Aegis Research Corporation

JAMES R. EASTER
Aegis Research Corporation

ABSTRACT

There has been a growing interest in using Cognitive Task Analysis (CTA) to understand the requirements of cognitive work and to provide a foundation for design of new decision-support systems. While CTA techniques have proved successful in illuminating the sources of cognitive complexity in a domain of practice and the basis of human expertise, the results of the CTA are often only weakly coupled to the design of support systems. A critical gap occurs at the transition from CTA analysis to system design, where insights gained from the CTA must be translated into design requirements. We describe an approach to bridging this gap. Our approach uses intermediate design artifacts to create a design thread that links the demands of the domain as revealed by the CTA, to the cognitive and collaborative processes that require support, through the elements of the decision aid that explicitly address those support requirements.

We present a visualization that was developed using this methodology to illustrate the approach. The visualization was developed to support military commanders in selecting among alternative courses of action.

5.1 INTRODUCTION

There is a tremendous need for powerful Decision Support Systems (DSS) to support humans in the increasingly complex work domains they face. The methodology presented in this chapter to "bridge the gap" between cognitive analysis and the design of effective decision aids is based on some fundamental premises chaining back from the fundamental goal of building effective DSSs—ones that "make the problem transparent to the user" (Simon, 1981). These premises constitute the underlying basis for the methodology described in this chapter. By understanding these underlying assumptions, the reader will have the necessary context to understand why the methodology evolved to its current state: an adaptation of fundamental research into a pragmatic engineering approach. The ultimate goal has been to develop DSSs that complement the human decision maker to form an integrated human-machine team capable of solving difficult, real-world problems more effectively than either individually.

Behind this goal is *Premise 1—Humans form a mental model of the domain as part of their understanding and problem solving*. In addition, they employ a variety of problem-solving strategies to reason on that mental model and make decisions, varying both by situation and by variations between human decision makers. These problem-solving strategies will be a situationally varying mix of sensory-motor responses (skill-based), actions based on stored rules and experience (rule-based), and behavior based on an internal representation of underlying, fundamental behavioral characteristics of the work domain (knowledge-based). Rouse and Morris' (1986) definition of mental models is functional in nature and will be the one utilized in this chapter. They define mental models as "the mechanism whereby humans are able to generate descriptions of system purpose and form, explanations of system functioning, observe system states, and prediction of future system states" (Rouse & Morris, 1986; p. 351). Experts employ "better" mental models—models with richer domain knowledge, more structure and interconnections, and a basis in the underlying principles of the work domain (Woods & Roth, 1988).

The requirement for the DSS to be "transparent" to the user, that is, to allow an understanding of the problem to flow into the human's decision making with virtually no cognitive effort results in *Premise 2—The DSS must itself embody a "knowledge model" of the domain that closely parallels mental models representative of expert human decision-making.* This is consistent with the notion of common frame of reference (Woods & Roth, 1988). The more closely the models parallel, the less cognitive effort it will take on the part of the human to "transform" the data of the DSS into the information needed to mentally solve the problem. As the degree of match approaches one, the concept of "transparent" becomes a reality. In addition to providing an effortless match for the expert operator, the knowledge model within the DSS can be viewed as an analytically derived *a priori* starting point for the novice user's own embryonic mental model of the domain. Also, the DSS's knowledge model can serve as a constant reminder of infrequently used concepts and knowledge even to experts.

To engineer such a knowledge model within the DSS results in *Premise 3—An effective DSS knowledge model is composed of functional nodes and relationships intrinsic to the work domain.* Explicit denotation of the cognitive tasks associated with nodes and relationships forms the basis of the support requirements of the DSS, which then must support each of the cognitive tasks across a variety of strategies. This DSS knowledge model therefore defines the basis for both the structure and content of the DSS itself; it becomes the overall design specification (Woods & Roth, 1988).

To deliver such a knowledge model requires *Premise 4—An adaptation of Rasmussen's abstraction hierarchy provides the needed representation of the abstract concepts and relationships applicable across all situations and strategies.* Functional nodes representing goals and functional processes to achieve those goals, at varying levels of abstraction, are a critical element differentiating expert (high-quality) problem solving from novice problem solving. In addition, the explicit relationships between those nodes provide the natural structure of the work domain. The functional abstraction hierarchy described in this paper is a pragmatic adaptation of that approach, but nonetheless based on the same premise: An abstract model of the functional nature of the domain provides the necessary exploration of highly abstract concepts of the domain while being sufficiently robust (independent of the physical particulars of the situation) to avoid brittleness in unexpected problem solving situations (Vicente, 1999). With this premise, the problem has come full circle—the domain defines the fundamental skeleton of the DSS's knowledge model to support a human dealing with that same domain.

These premises were organized as a reverse chain of logic from the starting requirement of "build a DSS to deliver improved human-machine decision making effectiveness," which exactly represents the evolution of the methodology described below. It provides a sense for the rationale underlying the various decisions that created the methodology. The remainder of the chapter is presented in a forward chain more representative of how the methodology is actually practiced for development of a particular DSS: When presented with a complex, real-world (naturalistic decision making) domain, how can a truly effective (transparent) DSS be developed?

To answer that question, there has been growing interest in using Cognitive Task Analysis (CTA) to understand the requirements of cognitive work and to provide a foundation for design of new decision-support systems (Schraagen, Chipman, & Shalin, 2000; Vicente, 1999). While CTA techniques have proved successful in illuminating the sources of cognitive complexity in a domain of practice and the basis of human expertise, the results of the CTA are often only weakly coupled to the design of support systems (Potter, Roth, Woods, & Elm, 2000). A critical gap occurs at the transition from CTA analysis to system design, where insights gained from the CTA must be translated into design requirements.

In this chapter we will describe our approach to bridging the gap between analysis and design. Our approach uses intermediate, decision-centered artifacts to create a bridge that links the demands of the domain as revealed by the cognitive analysis, to the cognitive processes (both individual and collaborative) that require support, through the elements of the decision aid that explicitly address those support requirements. Following the methodological discussion we will present an illustrative visualization that was developed using this methodology. The visualization was developed to support military commanders in selecting among alternative courses of action in applying combat power.

5.2 DEVELOPING A DESIGN THREAD FROM COGNITIVE ANALYSIS TO DECISION AIDING

Our approach is predicated on the premise that the design of advanced visualizations and decision-aids must explicitly reflect the fundamentals of the domain of practice and the demands it imposes on domain practitioners. We employ a structured, principled methodology to systematically transform the problem from an analysis of the demands of a domain to identifying visualizations and decision-aiding concepts that will provide effective support. The steps in this process include:

- Capturing the essential domain concepts and relationships that define the problem-space confronting the domain practitioners
- Identifying the cognitive demands/tasks/decisions that arise in the domain and require support
- Identifying the Information Requirements (IRs) to successfully execute these cognitive demands/tasks/decisions
- Defining the relationships between Decision Requirements (DRs) and user interface design concepts
- Exploring techniques to implement these design concepts into powerful visualization and decision support concepts.

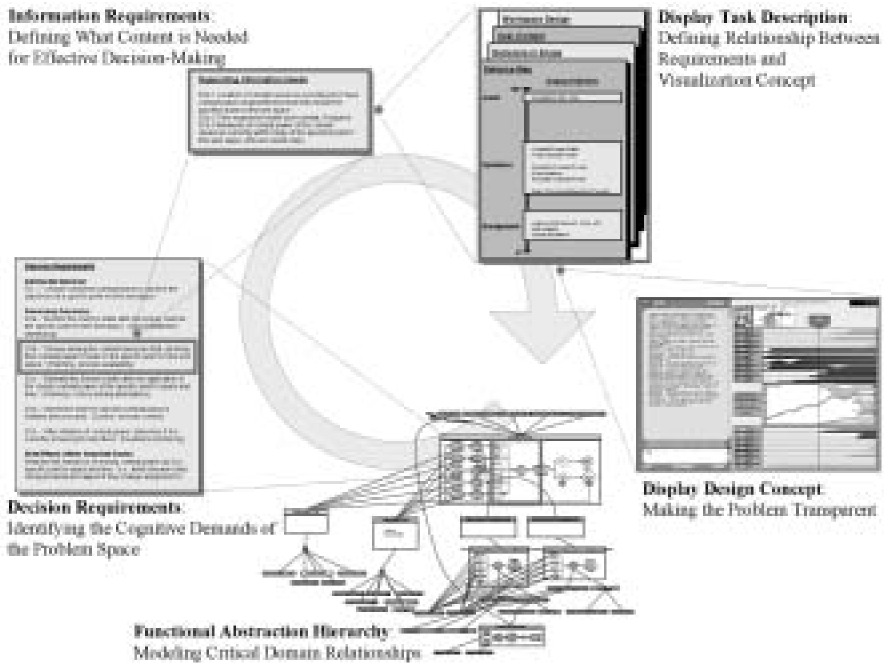

Figure 5.1: A sequence of analysis and design steps creates a continuous design thread that starts with a representation of domain concepts and relationships through development of decision support requirements to creation of visualization and aiding concepts and rapid prototypes with which to explore the design concepts.

In our analysis and design approach, we create design artifacts that capture the results of each of these intermediate stages in the design process. These design artifacts form a continuous design thread that provides a principled, traceable link from cognitive analysis to design. However, the spans of the bridge are constructed by the process of generating these artifacts, not the artifacts themselves. The artifacts serve as a mechanism to record the results of the process. Figure 5.1 provides a visual depiction of the sequence of methodological activities and associated design artifacts. The remainder of the paper describes and illustrates the approach.

5.2.1 Representing the Way the World Works—Building a Functional Abstraction Hierarchy

This methodology begins with a function-based goal-means decomposition of the system. This methodology has roots in the formal, analytic goal-means decomposition method pioneered by Rasmussen and his colleagues as a formalism for representing cognitive work domains as an abstraction hierarchy (e.g., Lind, 1993; Rasmussen, 1986; Rasmussen, Pejtersen, & Goodstein, 1994; Roth & Mumaw, 1995; Vicente, 1999; Vicente & Rasmussen, 1992; Woods & Hollnagel, 1987). A work domain analysis is conducted to understand and document the goals to be achieved in the domain and the means available for achieving them (Vicente, 1999). The objective of performing this analysis is to develop a structure that links the purpose(s) of individual controllable entities with the overall purpose of the system. This includes knowledge of the system's characteristics and the purposes or functions of the specific entities. The result of the first phase is a Functional Abstraction Hierarchy (FAH)—a multilevel recursive means-ends representation of the structure of the work domain—that anchors the first span of the bridge.

The work domain analysis is performed based on extensive interactions with expert practitioners in the domain and includes face-to-face interviews with the experts, watching the experts work in the domain, verbal protocol techniques, and other CTA and Cognitive Work Analysis (CWA) methods (see Potter, Roth, Woods, & Elm, 2000; Vicente, 1999). In practice, building an FAH is an iterative, progressively deepening process. It starts from an initial base of knowledge (often very limited) regarding the domain and how practitioners function within it. Then, complementary techniques are used to expand and enrich the base understanding and evolve a function-based model from which ideas for improved support can be generated. This process is highly opportunistic. Whether one starts by focusing on understanding the domain or by focusing on the knowledge and skills of domain practitioners depends on the specific local pragmatics. The key is to focus on progressively evolving and enriching the model so as to ultimately discover an understanding of the goal-driven characteristics of the domain that will lead to an understanding of the decisions practitioners are faced with in the domain.

The phrase "bootstrapping process" has been used to describe this process and emphasize the fact that the process builds on itself (Potter et al., 2000). Each step taken expands the base of knowledge providing opportunity to take the next step. Making progress on one line of inquiry (understanding one aspect of the field of practice) creates the room to make progress on another. For example, one might start by reading available documents that provide background on the field of practice (e.g., training manuals, procedures); the knowledge gained will raise new questions or hypotheses to pursue that can then be addressed in interviews with domain experts, and it will also provide the background for interpreting what the experts say. In turn, the results of interviews or exercises may point to complicating factors in the domain that need to be modeled in more detail in the FAH. This provides the necessary background to create scenarios to be used to observe practitioner performance under simulated conditions or to look for confirming example cases or interpret observations in naturalistic field studies.

The resulting FAH specifies the domain objectives and the functions that must be available and satisfied to achieve their goals. In turn, these functions may be abstract entities that need to have other, less abstract or less aggregated functions available and satisfied so that they might be achieved. This creates a decomposition network of objectives or purposes that are linked together from abstract goals to specific means to achieve these goals. For example, in the case of engineered systems, such as a process control plant, functional representations are developed that characterize the purposes for which the engineered system has been designed, and the means structurally available for achieving those objectives. In the case of military command and control systems, the functional representations characterize the functional capabilities of individual weapon systems, maneuvers, or forces and the higher-level goals related to military objectives.

The FAH provides a framework for making explicit the goals to be achieved in the domain and the alternative means available for achieving those goals. High-level goals, such as impacting a critical function, are decomposed into supporting lower-level subgoals. This provides the basis for identifying—through subsequent steps in the analysis and design process—the cognitive activities that arise in the domain and the information needed to support those decisions. The FAH enables the designer to determine where decision making is likely to be difficult due to the fundamental characteristics of the domain. For example, the FAH helps convey places in problem space where objectives compete with each other (e.g., where choices have to be made that require some level of sacrificing of one objective to achieve another, perhaps more heavily weighted, objective), or otherwise constrain each other (e.g., where the satisfaction of multiple goals needs to be considered in determining the best course of action).

Figure 5.2 depicts an example that illustrates these essential characteristics of an FAH:

- Processes may affect more than one goal—these side effects govern the operation of a process to achieve the goal of interest
- Each process can be modeled qualitatively to represent how it works to achieve a goal
- Relationships within the model can be recursive—processes can have requirements that are supported by more "abstract" process
- The term "hierarchy" is actually a misnomer; the structure of the model is actually a network
- Moving up through the network defines supported processes and impact on goal achievement; moving down defines supporting processes and requirements for goal achievement.

There are a growing number of examples of successful systems that have been developed based on a work domain analysis. Examples of functional

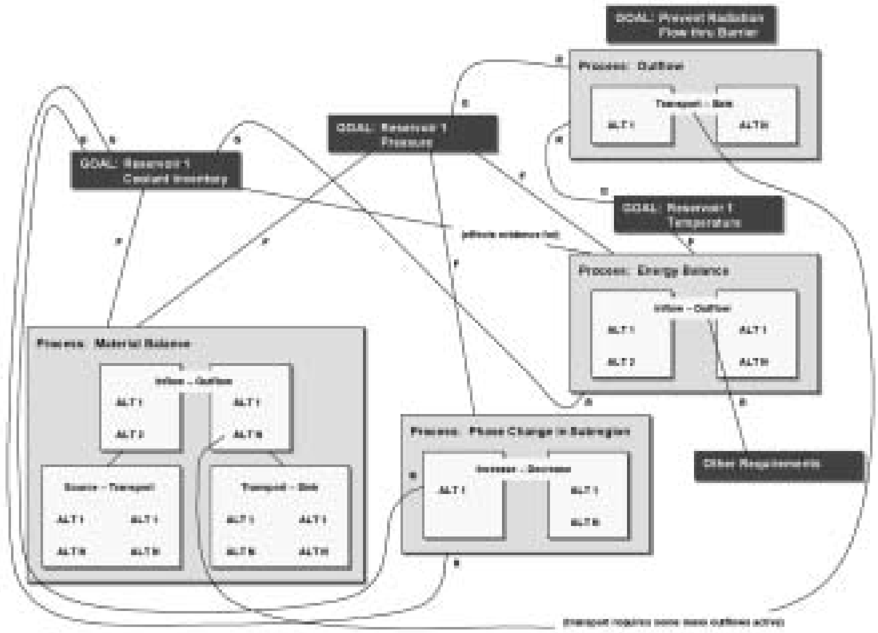

Figure 5.2: A sample goal-means decomposition from Woods and Hollnagel (1987) for one port (primary-system thermodynamics) of a nuclear power plant.

abstraction hierarchies and how they were used to design new visualizations and DSSs can be found in Roth, Lin, Kerch, Kenney, and Sugibayashi (in press) and Potter, Roth, Woods and Elm (2000). Examples of the application of this approach to model cognition and collaboration and to develop new online support systems in time pressured tasks such as situation assessment, anomaly response, supervisory control, and dynamic replanning include domains such as military intelligence analysis (Potter, McKee, & Elm, 1997), military aeromedical evacuation planning (Cook, Woods, Walters, & Christoffersen, 1996; Potter, Ball, & Elm, 1996), military command and control (Chalmers, Easter, & Potter, 2000), railroad dispatching (Roth, Malsch, Multer, Coplen, & Katz-Rhoads, 1998), and nuclear power plant emergencies (Roth, Lin, Thomas, Kerch, Kenney, & Sugibayachi, 1998).

5.2.2 Modeling Cognitive Demands—Deriving DRs

With the FAH representation of the work domain as the underlying framework, it is possible to derive the cognitive demands for achieving domain goals. In our methodology, we refer to these demands as DRs. Thus, the term "decision" is used in a broad sense. Based on the underlying premises of the modeling methodology, these decisions center around goal-directed behavior, such as monitoring for goal satisfaction and resource availability, planning and selection among alternative means to achieve goals, and controlling activities (initiating, tuning, and terminating) to achieve goals (Roth & Mumaw, 1995) as well as collaboration activities in team settings (Gualtieri, Roth, & Eggleston, 2000). By organizing the specification of operator DRs around nodes in the goal-means structure, rather than organizing requirements around predefined task sequences (as in traditional approaches to task analysis), the representation helps ensure that the resulting design concepts (on the system design side of the gap) reflect a decision-centered perspective. The resulting displays and decision-aids will thus support domain practitioners in understanding the goals to be achieved and what decisions and actions need to be taken to achieve these goals.

The cognitive demands that are derived from a cognitive analysis of the work domain constitute a second span in the bridge—DRs. DRs are tied directly to nodes in the FAH and provide an intermediate artifact that forms the essential part of the design thread, eventually providing an end-to-end connection from goal nodes in the FAH to supporting visualization and decision support concepts.

The FAH forms the basis for the structure of the decision-making activities that will be reflected in the DRs. For example, every goal node in the FAH has associated "goal monitoring" types of decisions. Likewise, processes have associated "process monitoring" decisions. Similarly, there will always be "feedback monitoring" types of decisions related to assessing whether actions are achieving desired results. Depending on the relationships between nodes in the FAH, there will be decisions related to prioritization of goals, selection of alternative means to achieve a particular goal, and monitoring side effects of actions.

An underlying "template" for this step in the analysis is presented in Table 5.1.

The key issue here is that this template is not meant to be a rote, "turn the crank" type of process. Rather, these questions are meant to be a guide to stimulate thinking about relevant decision-making in the context of a FAH model of the target work domain. Each domain is unique in the decision-making demands imposed on the human operators. As such, each work domain will require slightly different variants of these questions. Successful elucidation of DRs will also depend on corroboration from multiple data sources, including case studies, interviews, observations, etc. In addition, guiding insights can come from research on similar work domains as well as basic research on human cognition, decision-making, biases, and errors. For example, previous work on decision making in dynamic, high-risk worlds can guide analysis and interpretation of analogous worlds in terms of potential points of complexity, typical decision making difficulties and strategies, and critical characteristics of difficult problem-solving scenarios.

Table 5.1: Template of Typied Decision Types and Associated Knowledge Acquisition Questions. (Adapted from Roth & Mumaw, 1995.)

- Monitoring/Situation Awareness:
    - Goal Monitoring:
        - Goal Satisfaction: Are the function-related goals satisfied under current conditions?
        - Margin to Dissatisfaction: Are goal limits/restrictions being approached?
    - Process Monitoring:
        - Active processes: What processes are currently active? What is the relative contribution of each of the active processes to goal achievement? Are the processes performing correctly?
        - Process element monitoring: Are the individual processes and their components working as they are supposed to?
        - Automation monitoring: Are automated support systems functioning properly? What goals are the automated support systems attempting to achieve? Are these appropriate goals?
    - Feedback Monitoring:
        - Procedure adequacy: Is the current procedure achieving the desired goals?
        - Control action feedback: Are the operator control actions achieving their desired goals?
- Planning:
    - Goal priority: Which goal has the highest priority?
    - Process availability: What alternative processes are available for achieving the goals?
    - Choices among alternatives: Can an alternative process be deployed?
    - Consequences and side-effects of actions: What other processes or functions are affected by the current actions?
- Control:
    - Process control: How is the process controlled for process deployment, tuning for optimum performance, termination?
    - Manual take-over: If intervention is required, what actions must be taken?

5.2.3 Capturing the Means for Effective Decision-Making—Identifying Information Requirements

The next step in the process is to identify and document the information required for each decision to be made. Information requirements are defined as the set of information elements necessary for successful resolution of the asso-

1. This IR supports a DR of "monitor flow coefficient within operational limits."

ciated decision requirement. This set of information constitutes a third span in the bridge—IRs. The focus of this step in the methodology is on identifying the ideal and complete set of information for the associated decision-making.

IRs specify much more than specific data elements; it is data in context that becomes information (Woods, 1988, 1995). The data-to-information relationship can be complex and require a significant amount of computations and/or transformations. For example, in the case of a thermodynamic system, an IR might be "flow coefficient with respect to appropriate limits."[1] This requires the estimation of the parameter "flow coefficient" derived from model-based computations and sensor values and the comparison of that parameter against a limit referent. The degree of transformation required can vary from simple algebra to complex, intelligent algorithms. Potter et al. (1996) provide an example of IRs that were only able to be satisfied by an advanced planning algorithm and significant data transformations.

In addition, identifying IRs is focused on satisfying the DRs and is *not* limited by data availability in the current system. In cases where the required data is not directly available, this approach provides a rationale for obtaining that data (e.g., pulling data from a variety of previously "stove-piped" databases, adding additional sensors, or creating "synthetic" values). This is a critical change from the typical role that human factors engineers have had in the past (designing an interface after the instrumentation has been specified). Consequently, this type of an approach is fundamentally broader in scope than other approaches to interface design that do not consider the impact of IRs on system architecture specifications (Vicente, Christoffersen, & Hunter, 1996).

An interesting anecdote of this occurred in an interface design effort for a thermodynamic system (Potter et al., 1992). At this point in the process, the IR of "predicted liquid level in the accumulators versus current level over time" was identified to compensate for significant lags in the system in monitoring for system integrity. One of the engineers argued "but we don't have any way to sense 'predicted level'—that's our fundamental problem!" Slowly, another engineer in the room raised his hand and offered "my high-fidelity simulation of the system calculates that exact thing…I've just never talked about it because I didn't think it was of any value to anyone except me."

Just as the FAH representation provided the framework for the derivation of DRs, the DRs provide the essential context for the IRs because they indicate the factors (and thus information) that will need to be considered in making decisions. For example, in a "choice among alternative resources" type of decision, the choice requires information about the availability of the alternatives (supporting relationships), current tasking of those alternatives and the impact of selecting it for the task under consideration (side-effects), and specific performance capabilities of the alternatives (lower level functional properties of the alternatives).

5.2.4 Linking DRs to Aiding Concepts—Developing and Documenting a "Model of Support"

Once the FAH has been augmented with critical decision and information requirements, it becomes a solid foundation for the development of aiding concepts. The objective is to design visualization and decision support concepts to reflect the IRs that, as a result of the linkage back to the FAH, are organized into a virtual "information space" explicitly replicating the domain structure captured in the FAH. To accomplish this objective, this task develops the mapping between information on the state and behavior of the domain (i.e., decision and information requirements uncovered) and the syntax and dynamics of the visualization or decision aid being developed. From an interface design perspective, the goal is to reveal the critical IRs and constraints of the decision task through the user interface in such a way as to capitalize on the characteristics of human perception and cognition. This approach is consistent with cognitive engineering principles that have variously been called *Representational Aiding* (Bennett & Flach, 1992; Roth, Malin, & Schreckenghost, 1997, Woods, 1995) and *Ecological Interface Design* (Vicente & Rasmussen, 1990, 1992; Reising & Sanderson, 1998).

The display concept and how it supports the cognitive tasks is then captured in a Display Task Description (DTD). The DTD defines the goals and scope of a display in terms of the cognitive tasks it is intended to support (and thus a defined target region of the FAH). It also provides a specification of the supporting information and graphic elements required to support the cognitive tasks. A DTD is another span of the bridge that helps to link the decisions within the work domain to the visualization and decision support concepts intended to support those decisions. In many cases, multiple design concepts may be generated that support a given set of decisions. These alternative solutions can be captured in a DTD.

An advantage of a DTD is that it requires designers to be more explicit about the specific cognitive activities that a given visualization or decision support concept addresses. DTDs specify the decisions to be supported and the cognitive performance objectives that the display or decision aid is intended to achieve and the information (not just data) that must be conveyed. Explicit links are made between particular aspects of the display concepts and specific cognitive demands they are intended to support. As such they constitute explicit hypotheses—a model of support—that can be empirically evaluated. As a consequence, DTDs enable more informed and pointed testing of the effectiveness of the proposed aiding concepts.

Another advantage of DTDs is that they enable designers and evaluators of designs to clearly distinguish and independently evaluate the objectives of a display in terms of intended support from the particulars and aesthetics of its implementation. One can ask "are the support objectives of this display correct and complete?" as well as "does a particular embodiment of the display concept achieve the intended support objectives?" This is a key aspect of bridging the gap between analysis and design.

The DTD also represents a configuration management tool, critical for ensuring coverage of the functional decision space across all displays and display elements. The DTD represents a shift in focus from "what" is to be displayed to "how," including annotations on relative importance that maps to relative salience on the visualization, etc. The DTD is not only a compilation of information developed earlier, it has the added value of a more complete description of the behaviors and features needed to communicate the information effectively as well as an allocation of the IRs across the entire set of displays within the workspace. When done correctly it is still in the form of a "requirement" and not an implementation. This artifact becomes a key transition between the cognitive system engineer, the system developer, and the system tester.

5.2.5 Developing Prototypes That Instantiate the Aiding Concept

The introduction of new technology inevitably transforms the work domain and the demands placed on domain practitioners, often in unanticipated ways. New technology introduces new error forms; new representations change the cognitive activities needed to accomplish tasks and enable the development of new strategies; new technology creates new tasks and roles for people at different levels of a system. Changing systems change what it means for someone to be an expert and change the kinds of errors that will occur.

Given this transformation, developers face the "envisioned world" problem of the unforeseen impacts of the introduction of new technology (Dekker & Woods, 1997; Smith, Woods, McCoy, Billings, & Sarter, in press; Woods, 1998). A similar phenomenon has been noted in the computer-human interaction literature where it is referred to as the "task-artifact cycle" (Carroll, Kellogg, & Rosson, 1991). Concepts for new visualizations and DSSs (aiding concepts) represent hypotheses about what will provide effective support to domain practitioners in the envisioned world. Rapid prototypes of aiding concepts that implement the DTDs become tools for discovery (Potter, Roth, Woods, & Elm, 2000). By exploring the impact of prototypes that embody aiding concepts, it becomes possible to evaluate the effectiveness of the envisioned support systems (i.e., test the hypotheses) as well as to identify additional support requirements and unanticipated consequences of the introduction of the new technologies.

The envisioned world problem means that system developers must face a challenge of prediction:

- What will be the cognitive demands of the envisioned world?
- How will the envisioned support concepts shape cognition

and collaboration?

- How will practitioners adapt artifacts to meet their own goals, given mismatches to the actual demands and the pressures they experience?
- How will the new technology impact a domain that doesn't yet exist or is evolving?

The goal of such predictions is to influence the development process so that new decision aids provide effective, robust decision support. This can be accomplished by developing prototypes of the visualization and decision support concepts as specified in the DTD. These prototypes provide a concrete instantiation of the aiding concepts specified in the DTD. They can be used to explore the viability of the aiding concept. Then, each opportunity to assess the utility of the prototype can also provide additional understanding of the requirements for effective support. Thus, these assessments can serve to enrich and refine the initial FAH and identify additional decision and information requirements that were missed in the original analytic process. Note that extending the analysis to encompass exploration of the envisioned world in a closed-loop, iterative manner contrasts with the narrow view of cognitive analysis as an initial, self-contained technique whose product is handed-off to system designers in a waterfall model approach.

The methodology and resulting design artifacts described in these sections provide the blueprint for the development of DSSs designed around the fundamental, underlying demands of the work domain to deliver significantly improved human-machine decision-making effectiveness. The remainder of this chapter will provide an instantiation of these artifacts for a component of a DSS prototype designed to support military commanders in selecting among alternative courses of action in applying combat power.

5.3 AN ILLUSTRATIVE CASE: AN ECOLOGICAL INTERFACE SUPPORTING COMMAND DECISION MAKING

A decision aid that was recently developed to support military command decision-making can be used to illustrate the breakthroughs achieved from this methodology. The display was developed as part of the DARPA Command Post of the Future program, as a means of illustrating the use of this approach to support the development of powerful new visualizations (Logica Carnegie Group, 2000a). The discussion of this case study will follow the intermediate artifacts that provide a design thread from cognitive analysis through prototype implementation.

2. The electronic tactical decision game (eTDG) was developed by John Schmidt as part of the Command Post of the Future program.

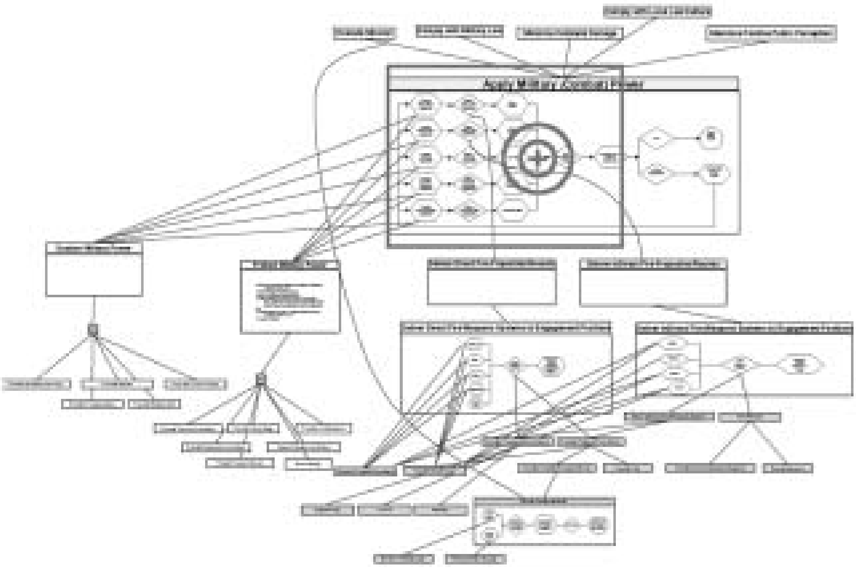

Figure 5.3: Functional abstraction hierarchy of military command and control with the "Apply Military (Combat) Power" portion highlighted.

5.3.1 Functional Abstraction Hierarchy (FAH)

As part of the program an FAH was developed for military command and control. The FAH is presented in Figure 5.3. It depicts the goals, functional processes necessary to achieve the goals, and the subgoals that result from the need to support the functional processes.

A variety of knowledge elicitation techniques were used to bootstrap an initial understanding of the decision problems faced by military commanders, how military commanders conceptualize the problem space, and the complications that arise in the domain that increase the difficulty of decisions.

Knowledge acquisition activities included:

- Reviewing military documents
- Conducting structured interviews with military commanders at different levels in the command chain, including recently retired general officers with combat experience
- Participating in electronic tactical decision games that were conducted as part of the Command Post of the Future program and included recently retired general officers as game participants[2]
- Presenting draft versions of the FAH to military commanders and revising them based on their feedback. This was an iterative process that occurred over several cycles.

One of the findings of the knowledge elicitation process is that military commanders think about military units in terms of the abstract concept "combat power" that military units possess and can generate. This abstract concept includes not only the number and type of military equipment that the unit possesses, but also less tangible factors such as the unit's morale, fatigue, etc., as well as a number of external factors including terrain. In Figure 5.3, "Apply Combat Power," has been generalized to "Apply Military (Combat) Power" to reflect the fact that the modern day military is often asked to engage in non-combat operations (e.g., providing humanitarian aid).

The function of applying combat power (or more generally military power) is performed in the context of meeting several higher-level goals. The most direct higher-level goal is to satisfy mission objectives. There are other higher-level goals, however, that must be taken into account. These include "complying with military law," "minimizing collateral damage," "complying with local laws and cultures," and a need to attain "positive public perception" of the operation. These additional goals can place constraints on achieving mission goals. The fact that high-level goals can compete with each other is reflected in the FAH.

The military command and control FAH illustrates one of the important benefits of developing a functional abstraction hierarchy representation of the domain. The exercise of developing an FAH enables the cognitive analyst to see beyond the physical level of description of the domain and to begin to understand and represent the domain at higher levels of abstraction. The concept of "combat power" that is represented in the military command and control FAH is a good example. Combat power can be thought of as a "commodity." The commander is given resources (troops, planes, tanks, logistics support) that have many complex functional properties and interdependencies that affect the amount of combat power that can be delivered at a given point in space and time. As such, the commander can be thought of as the manager of a very precious commodity, seeking ways to maximize the combat power that can be brought to bear to achieve a particular mission.

Military commanders routinely think in terms of abstract concepts such as "relative combat power" in making decisions about movement of troops and equipment to achieve a mission objective. Yet the current tools available to them (e.g., physical maps of the terrain with icons representing placement of troops and equipment) provide a much more physical representation of the problem space, making the estimation of relative combat power difficult and prone to error. Anecdotally, the IR for this domain included "relative combat power over time" a very untraditional concept that created significant excitement since it solidified some doctrinal writings about controlling the "tempo" of the conflict.

A second important attribute illustrated by the military command and control FAH is the value of the conceptual structure it provides. The focus of effort in developing a FAH is on constructing a representation that captures the important concepts and interrelationships in the domain at appropriate levels of abstraction and makes the goals, means, relationships, and constraints explicit. Initially, the emphasis is placed on shaping the representation around key concepts (goals and processes) and relationships. Once the right structure is in place, supporting details (i.e., less abstract concepts such as procedures, and functional capabilities) can be added to the FAH to flesh out the model.

It has often been our experience that when we finally get the representation "right" (after many iterations) the resulting FAH looks simple and is readily understood and accepted by domain Subject-Matter Experts (SMEs). We have coined the term "of course test" to describe the typical reaction of domain practitioners to an FAH that successfully crystallizes the most important abstract concepts and inter-relationships in the domain. It should be pointed out, however, that while the FAH may look simple and obviously true to the SMEs once it is presented to them, the conceptual structure represented in the FAH is not something that the domain experts could have spontaneously generated on their own. Often the concepts represented in the FAH are things that the domain experts understand implicitly, and readily resonate to once shown, but could not generate themselves.

The power of this abstraction hierarchy is that it allowed the cognitive analysts to recognize the need to create displays and decision-aids that allow the commander to visualize and control the domain at this higher, more abstract level of goal achievement. As part of the program a display was created that illustrates this point.

The objective of the display was to support commanders in choosing the appropriate military power to achieve mission objectives. The display was targeted at combat applications where the mission objective is to engage and defeat an enemy. In this combat context the probability of meeting mission objectives is a function of the ratio of friendly to enemy combat power at the point in time and space where the engagement is to take place.

The "choose combat power" display was designed to support commanders in deciding: (1) at what location to engage the enemy; (2) when to engage the enemy; and (3) what combat resources to deploy to achieve an acceptable potential to defeat the enemy. While this paper focuses on the "choose combat power" display, the display is intended to be only one of a suite of displays that would encompass a complete battle command decision-aid.

The FAH, shown in Figure 5.3, provided the starting point for developing the visualization to support commanders in choosing the appropriate combat power to carry out the mission objective(s). The starting point is the "apply military (combat) power" node in the FAH. This entails choosing the combination of resources that the commander believes will effectively carry out the mission.

Table 5.2: DRs Associated With the "Choose Combat Power" Visualization Concept

DRs
DR1a—"Determine the point in time when the enemy will reach a specified point in space and will monitor the enemy's combat power over time." (*This decision is driven by the need to establish the mission requirements (i.e., evaluate enemy combat power) so as to understand the friendly combat power that will be required to achieve mission goals (i.e., defeat the enemy)).*
DR1b—"Choose among the friendly combat resources that can bring their combat power to bear at the specific point in space and time." (*This decision considers the availability of resources to achieve mission goals.)*
DR1c—"Estimate the potential to defeat the enemy after the application of the chosen power at the specific point in time and space." (*This decision is related to the choice among the various alternatives.)*
DR1d—"Determine the impact of moving combat power to a specific point in space and time, i.e., what else were these resources committed to and what will happen if they change their assignments/committments?" (*This decision considers the side effects of choosing various alternatives.)*

5.3.2 Decision Requirements

The next step in the analysis and design process was to utilize the structure of the problem space map (i.e., the FAH) and derive the supporting decisions for accomplishing the objective in question. In this specific context, the "apply military (combat) power" node in the FAH suggested the decision of "choose combined combat power to achieve the objectives at a specific point in time and space." The decision is written in a generic form (i.e., independent of the particular battle or terrain that originally exposed the need for the decision). This is part of the process of making the DSS not situationally dependent. Then, the specific set of DRs were derived in part by imposing the template of generic questions (see Roth & Mumaw, 1995). The set of DRs that was derived through this process is shown in Table 5.2.

5.3.3 Supporting Information Requirements

Associated with each of these decisions is requisite domain-specific information needed to inform the decision. A description, in generic terms, of the type of information necessary to support the "choose combat power" decision is given in Table 5.3. The identification numbers for IRs cross-reference the DRs they support.

Table 5.3: Supporting Information Requirements Associated With the "Choose Combat Power" Decision Requirement

Information Requirements
DR 1a:
IR 1a.1 – "Expected arrival time of enemy combat resources at the specified point in space," (i.e., the lead unit, as well as other follow-on units).
IR 1a.2 – "Estimated measure of combined enemy combat power at the specified point in space, beginning at the arrival time of the first enemy unit and extending through follow-on units."
DR 1b:
IR 1b.1 – "The time required for selected friendly combat resources to reach the specified point in space."
IR 1b.2 – "Estimated measure of combined combat power of the selected friendly combat power resources once they reach the specified point in space."
DR 1c:
IR 1c.1 – "Measure of combat power ratio of friendly to enemy combat power beginning with the arrival of the first unit (friendly or enemy) over time."
IR 1c.2 – "Indication of combat power ratios required to defeat the enemy under different battle conditions (i.e., doctrinal / procedural referent information)."
IR 1c.3 – "Location of alternative resources of both friendly and enemy combat power that could be brought to bear."
IR 1c.4 – "The time required to bring to bear the combat power of these alternative friendly and enemy combat resources."
IR 1c.5 – "Measures of cumulative combat power of both friendly and enemy resources as additional friendly and enemy resources are selected (over a specified window in time)."

5.3.4 Display Task Description

Based upon the FAH, the DRs, and the IRs, a DTD was created to be used to initiate the visualization design. As mentioned earlier, the DTD provides a specification of the specific supporting IRs for the resulting graphic elements to support the cognitive tasks/decisions. If the graphic elements and visualizations effectively convey the IRs identified during the IRs analysis, we have successfully bridged the gap between the essential demands of the work domain and the resulting decision support concepts. In most cases, this also requires consideration of allocation of functional/decision/information scope for each display. However, for the present example, we have not included that broader scope issue in this discussion. Table 5.4 presents the display task description for the "choose combat power" display.

Table 5.4: Display Task Description for the "Choose Combat Power" Display

Display Task Description	
Context	
This display is intended for combat applications where the mission objective is to engage and defeat an enemy. It should assist the commander in choosing combined combat power and managing the space/time tradeoff to control the battlespace. Specifically it is to assist in deciding: (1) at what location to engage the enemy; (2) when to engage the enemy; and (3) what combat resources to deploy to maximize the potential to defeat the enemy.	
Decision and Information Requirements	**Visualization Needs**
DR 1a—"Determine the point in time when the enemy will reach a specified point in space and will monitor the enemy's combat power over time." **IR 1a.1—**"Expected arrival time of enemy combat resources at the specified point in space." (i.e., the lead unit, as well as other follow-on units). **IR 1a.2—**"Estimated measure of combined enemy combat power at the specified point in space, beginning at the arrival time of the first enemy unit and extending through follow-on units."	Visualize the time required for enemy unit(s) to reach location designated by commander, and visualize cumulative enemy combat power at that location as a function of time.
DR 1b—"Chose among the friendly coombat resources that can bring their combat power to bear at the specific point in space and time." **IR 1b.1—**"The time required for selected friendly combat resources to reach the specified point in space." **IR 1b.2—**"Estimated measure of combined combat power of the selected friendly combat power resources once they reach the specified point in space."	Visualize the time required for friendly "combat power" unit to bring their combat power to bear on the designated point (in space and time) and visualize cumulative friendly combat power at that location as a function of time.
DR 1c—"Estimate the potential to defeat the enemy after the application of the chosen power at the specific point in time and space." **IR 1c.1—**"Measure of combat power ratio of friendly to enemy combat power beginning with the arrival of the first unit (friendly or enemy) over time." **IR 1c.2—**"Indication of combat power ratios required to defeat the enemy under different battle conditions (i.e., doctrinal/procedural referent information)." **IR 1c.3—**"Location of alternative resources of both friendly and enemy combat power that could be brought to bear." **IR 1c.4—**"The time required to bring to bear the combat power of these alternative friendly and enemy combat resources." **IR 1c.5—**"Measures of cumulative combat power of both friendly and enemy resources as additional friendly and enemy resources are selected (over a specified window in time)."	Visualize relative friendly to enemy combat power ratio at the designated point in space as a function of time, and compare to combat power ratio required to defeat enemy under different battle conditions. Visualize uncertainty of estimates of Combat Power Ratio. Visualize changes in enemy and friendly combat power and combat power ratio as selected enemy and/or friendly units are added or removed.

5.3.5 Rapid Prototype of "Choose Combat Power" Display

A rapid prototype of the "choose combat power" display was developed to provide a concrete instantiation of the concepts embodied in the DTD. Two versions of the rapid prototype were developed. First, a static "storyboard" version of the display was built to serve as a vehicle for communicating the display concepts. Later, a software prototype was developed that is called Joint Operations Environment (JOE) to explore the dynamic behavior when driven from a simulated dataset.

Figure 5.4 presents a "scenario" that illustrates the context in which a visualization such as JOE would be used. The commander would begin by viewing a physical representation of the terrain with enemy and friendly units and their positions identified. This is similar to the types of geographic displays that

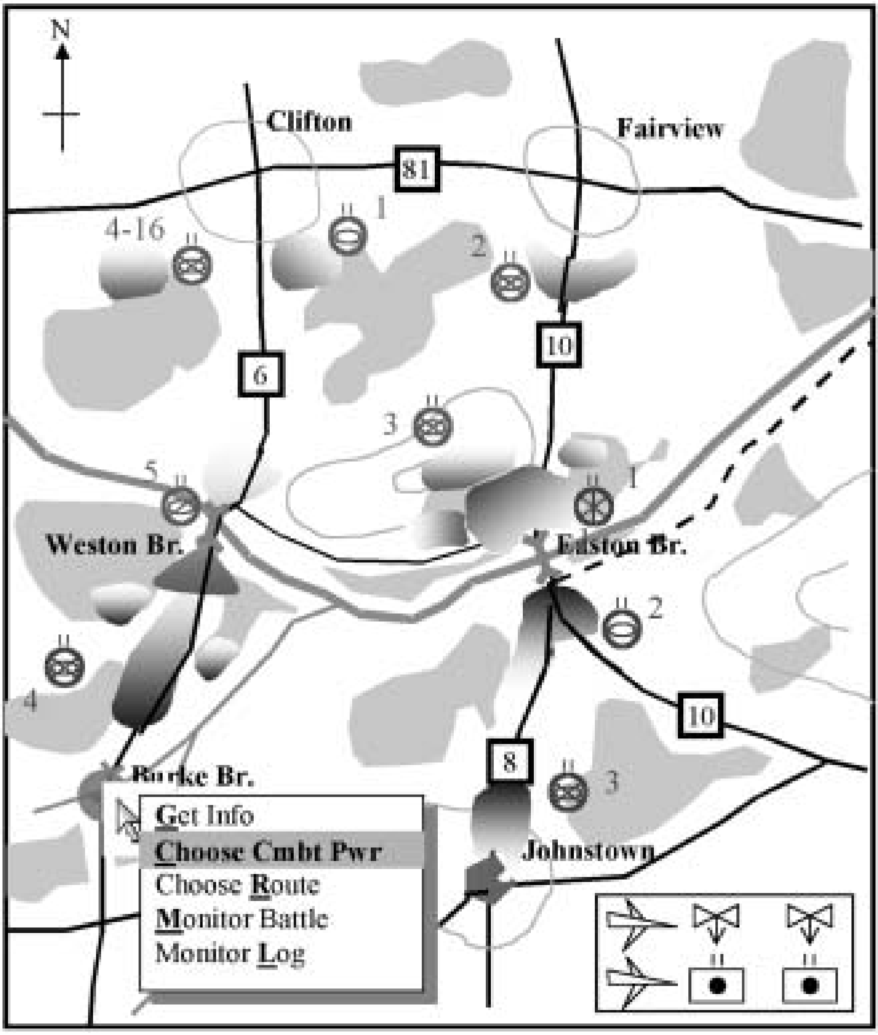

Figure 5.4: Geographic map providing the scenario context for the "choose combat power" visualization. This is classically the DSS a commander uses (along with some outboard math) to make the decision.

commanders use today. This scenario starts with the "Red" (enemy) forces attempting to deny the "Blue" (friendly) forces bridgehead at Weston and Easton bridge. Initially, Blue had decided to use the 2nd and 3rd battalions to support the advantage gained at Easton bridge. This would allow Blue to capture Fairview and possibly Clifton. However, at that point, Red's 5th battalion breaks through Blue's position at Weston bridge and threatens a flanking

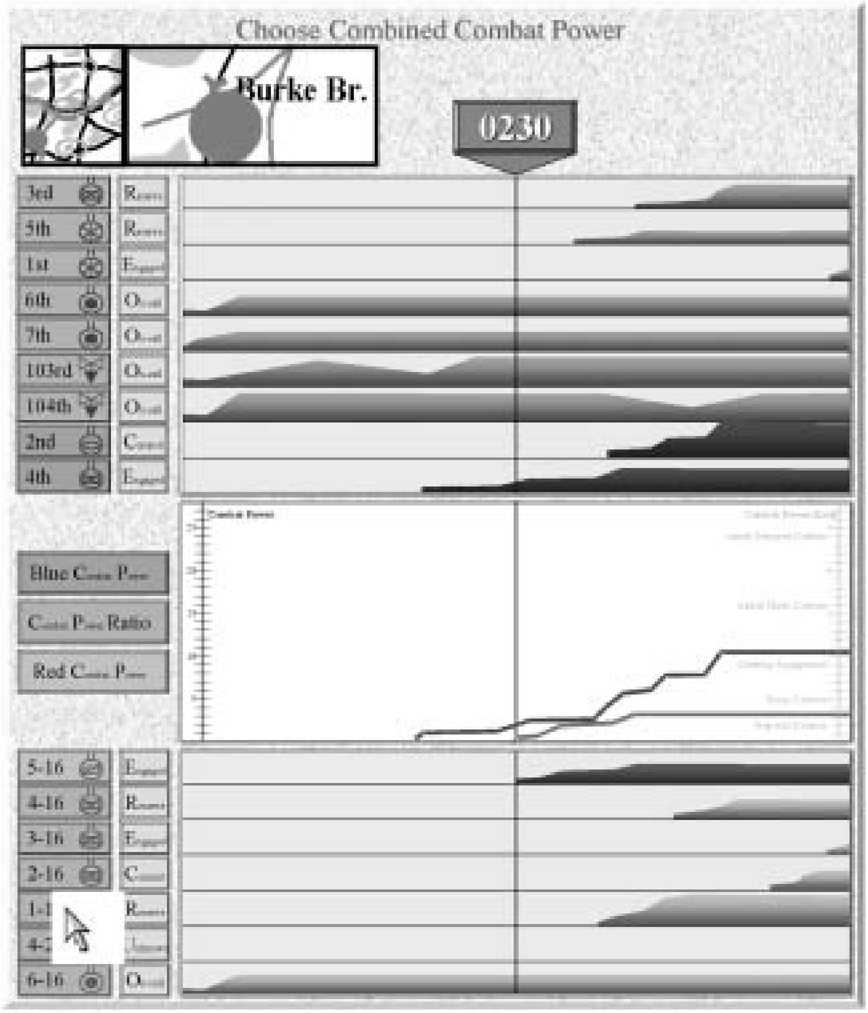

Figure 5.5: JOE visualization showing combat power of Red and Blue forces at the time that the 2nd and 4th Blue battalions and 5th Red battalion reach the "choke point" south of the Burke bridge.

attack from the west. At this point the Blue commander recognizes the need to interdict the Red 5th battalion that is pursuing the Blue 4th battalion. The commander considers using the 2nd Armor to interdict the Red 5th battalion at the "choke point" south of the Burke bridge.

To evaluate whether the 2nd Armor is sufficient to achieve the mission goal, the commander brings up the "choose combat power" display. The commander selects the choke point south of the Burke bridge as the designated location

Table 5.5: Mapping of Graphical Elements to DRs for the "Choose Combat Power" Display

Decision and Information Requirements	Visualization Needs	Graphical Elements
DR 1a – "Determine the point in time when the enemy will reach a specified point in space, and monitor the enemy's combat power over time." **IR 1a.1** – "Expected arrival time of enemy combat resources at the specified point in space." (i.e., the lead unit, as well as other follow-on units) **IR 1a.2** – "Estimated measure of combined enemy combat power at the specified point in space, beginning at the arrival time of the first enemy unit and extending through follow-on units."	Visualize the time required for enemy unit(s) to reach location designated by commander, and visualize cumulative enemy combat power at that location as a function of time.	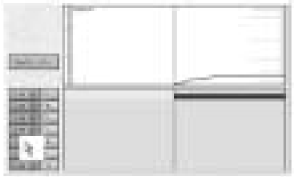
DR 1b – "Choose among the friendly combat resources that can bring their combat power to bear at the specific point in space and time." **IR 1b.1** – "The time required for selected friendly combat resources to reach the specified point in space." **IR 1b.2** – "Estimated measure of combined combat power of the selected friendly combat power resources once they reach the specified point in space."	Visualize the time required for friendly "combat power" unit to bring their combat power to bear on the designated point (in space and time) and visualize cumulative friendly combat power at that location as a function of time.	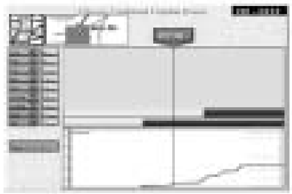
DR 1c – "Estimate the potential to defeat the enemy after the application of the chosen power at the specific point in time and space." **IR 1c.1** – "Measure of combat power ratio of friendly to enemy combat power beginning with the arrival of the first unit (friendly or enemy) over time." **IR 1c.2** – "Indication of combat power ratios required to defeat the enemy under different battle conditions (i.e., doctrinal / procedural referent information)." **IR 1c.3** – "Location of alternative resources of both friendly and enemy combat power that could be brought to bear." **IR 1c.4** – "The time required to bring to bear the combat power of these alternative friendly and enemy combat resources." **IR 1c.5** – "Measures of cumulative combat power of both friendly and enemy resources as additional friendly and enemy resources are selected (over a specified window in time)."	Visualize relative friendly-to-enemy combat power ratio at the designated point in space as a function of time and compare to combat power ratio required to defeat enemy under different battle conditions. Visualize uncertainty of estimates of Combat Power Ratio. Visualize changes in enemy and friendly combat power and combat power ratio as selected enemy and/or friendly units are added or removed.	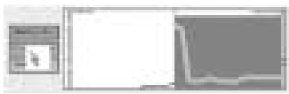

of the engagement and the 2nd and 4th Blue battalions and the 5th Red battalion as the units involved in the engagement. JOE computes the time it will take for these units to reach this designated location and the combat power that the Red and Blue sides will be able to bring to bear at that location over time.

Figure 5.5 presents the resulting JOE display. The display has three major areas. The top third of the display shows all the Blue units that could potentially be deployed, their current status, the time at which they would reach the designated location, and the combat power they could bring to bear over time. The shaded areas represent amount of combat power as a function of time. The horizontal axis represents time, with the left most representing the current time. The heights of the shaded areas represent the amount of combat power

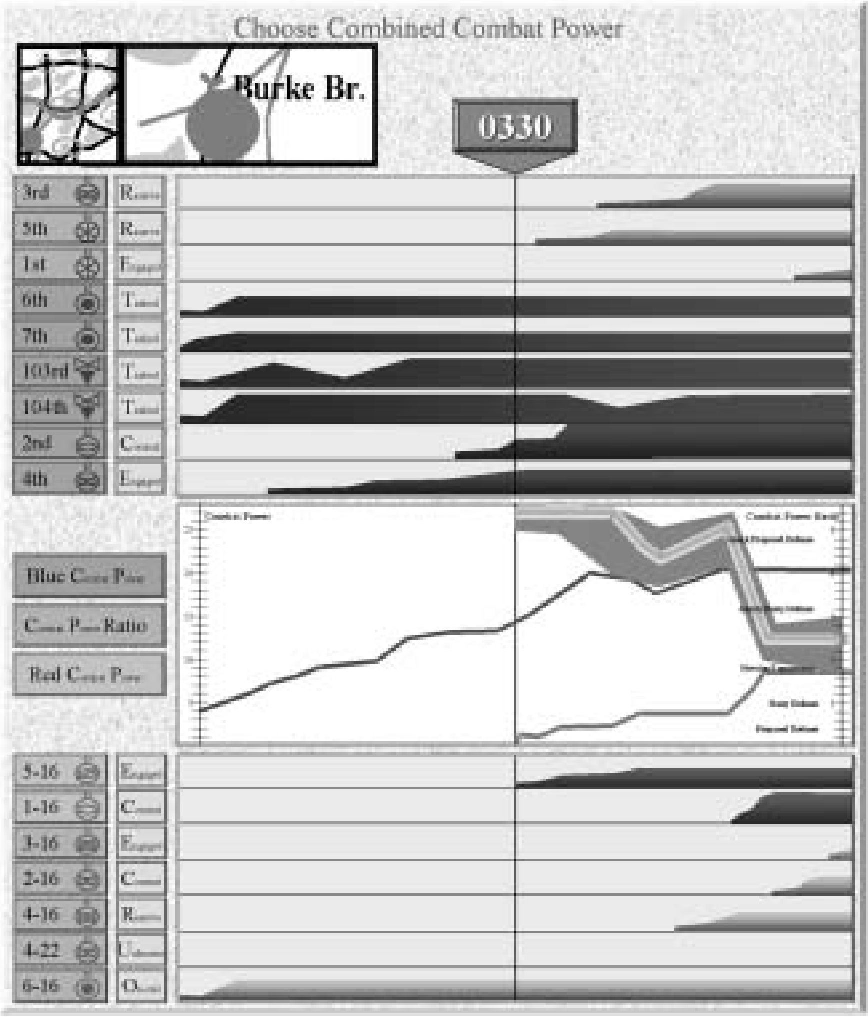

Figure 5.6: Updated JOE display reflecting the changes in relative combat power based on additional units selected. The thick line with bands around it illustrates a feature in JOE to display estimated combat power ratio (Blue/Red with bands of uncertainty around it).

at a given point in time. The units that are highlighted (appear darker) are the units that the commander has selected. Similarly the bottom third of the display shows the Red units and their combat power, with the highlighted unit (5th Battalion) indicating the Red unit selected by the commander. The middle third of the display is used to present the cumulative combat power of the combined Red and (separately) combined Blue units selected as a function of time. Thus, the higher curve in the center area represents the combined combat power of the 2nd and 4th Blue battalions over time. Combat power increases significantly when the 2nd Battalion reaches the designated location. The vertical line indicates the estimated time of arrival of the first Red unit. These three major areas map directly to the information and DRs contained within the DTD, as indicated in Table 5.5.

The JOE display can be used to visually compare the combat power of the Blue and Red units at the time that the Red unit first reaches the designated location. The graphic presentation of combat power as a function of time makes it visually apparent that the Red 5th battalion will reach the designated point before the Blue 2nd battalion can get there. As a result, the difference in combat power between the Blue and Red units at the time of the initial engagement will be small. Therefore the goal of defeating the Red units is unlikely to be achieved.

Figure 5.6 illustrates two additional features of JOE. First, JOE can be used to explore alternative combat resource choices by adding or removing enemy and friendly units to assess the impact it would have on combat power at a given point in time and space. Second, Joe can be used to display combat power ratio of “friendly” to “enemy” forces directly as well as bands of uncertainty around the estimate.

Figure 5.6 illustrates the change to the display when the commander decides that the Red first armor battalion is also likely to join the engagement. The commander decides that additional Blue resources (artillery and helos) will need to be brought in. The commander selects these additional resources and the JOE display is updated to reflect the changes in Red and Blue combat power that occur as a result of introducing these additional units. The commander also requests that combat power ratio be visually displayed. The thicker line in the center of Figure 5.6, an estimate of combat power ratio (Blue/Red) also represents the uncertainty band around the estimate of combat power ratio. A combat power ratio scale appears on the right with indications of combat power ratios recommended for different types of engagements based on conventional military guidance. As JOE makes visually apparent, in the new set of choices, there is enough Blue combat power brought to bear to make it highly probable that the Blue forces will achieve their mission.

As mentioned earlier, while this paper has focused on the “choose combat power” display, the display is intended to illustrate one of a suite of displays that would encompass a complete battle command DSS. As is suggested by the menu that appears in Figure 5.3, additional displays would address other aspects of battle command such as selection of routes and monitoring of the battle.

In addition, display navigation mechanisms would be provided to enable the commander to more fully understand the basis for the combat power values. For example the commander should be able to view the factors and values that contributed to combat power calculations and make changes to the values and factor weightings as judged appropriate. This is consistent with principles for design of effective decision aids that include the importance of making the basis of recommendations transparent to the user (the principle of decomposability) and the importance of enabling the user to direct the decision aid (the principle of directability) described in Roth, Malin, and Schreckenghost (1997).

Mechanisms would also be provided to enable the commander to assess the impact of selecting particular units to join an engagement on the achievement of other goals. Specifically, the commander should be able to select a unit being considered for a particular engagement, be shown the current (or other planned) commitments for this unit, and the potential impact on achieving those goals if the unit were to be reassigned. This supports the decision requirement to be able to assess the side effects of decisions, identified in Table 5.2.

5.4 DISCUSSION

The "choose combat power" visualization and the JOE rapid prototype of the display concept provide a concrete illustration of how a structured, principled methodology can systematically transform the problem from an analysis of the demands of a domain to identifying effective decision-aiding concepts. The process of generating intermediate design artifacts can be used to build a traceable bridge from cognitive analysis to design. The process of developing the JOE display concept began with cognitive analysis of the domain. A functional abstraction hierarchy was developed that captured the goals in the domain, the means available for achieving them, and the goal constraints and interactions inherent in the domain. The abstract concept of "combat power" and the central role it plays in choosing resources to bring to bear and accomplish a military mission emerged out of this analysis. A second insight that emerged out of the analysis was the importance of representing combat power as a function of time. It takes time for resources to be assembled and moved to a specified location. One of the breakthrough insights in developing the "choose combat power" display was the realization that such a visualization would allow commanders to explicitly manage the time of the culminating point of the engagement. These insights provided the basis for development of the "choose combat power" display concept and the JOE prototype. JOE enables domain practitioners to visualize "combat power" as a function of time and to manipulate it directly. It provides a clear example of an "ecological" interface in that it translates an abstract, functional concept (relative combat power as a function of time), into a concrete visualization that can be apprehended perceptually.

The "choose combat power" case study illustrates the value of performing a multilayered, work domain analysis based design process for generating powerful visualization concepts. First, the process of developing the FAH representation was central to the identification of the key abstract concept. This concept would not have been recognized without developing the FAH representation, as evidenced by typical physical map-based displays in widespread existence. Second, the process of defining the DRs, IRs, and display task descriptions (intermediate artifacts created to establish a traceable link between the analysis of domain demands and the display concept that emerged) serve as critical thought/design steps in the process. They force the analyst to revisit the FAH and consider issues possibly overlooked, be explicit about the cognitive tasks associated with the nodes in the model, and consider the often complex transformation between data and information. As a whole, they successfully serve to "bridge the gap" between the underlying cognitive analysis and resulting DSS design.

In this paper, the steps in the cognitive systems engineering process are presented as if they are performed in a strictly sequential order. First domain analysis, then DRs analysis, then IRs analysis, etc. It is presented this way for expository simplicity. In practice the process is much more parallel, opportunistic, and iterative in nature than presented here. For example, it is not unusual for an initial visualization idea to emerge before the complete DRs it is intended to support or the supporting information needs have been clearly articulated. The order in which the artifacts are produced is not as important as the fact that all artifacts are eventually produced that provide a functional description of the cognitive and decision tasks that the display is intended to support, and the information and display elements that provide the required support. As mentioned previously, the multiple spans of the bridge are constructed by the process of generating these artifacts, not the artifacts themselves. The key point is that it is the process of generating these artifacts that forces the cognitive analysts to think about the problem in a systematically transformed manner and capture the evolving requirements in a manageable sequence of steps.

The generation of intermediate artifacts that model the structure of the work domain, the demands and cognitive activities to be supported, the IRs for these cognitive activities, and decision support concepts designed to provide this support are needed to provide a traceable link from analysis, to design requirement, to display concept. The approach outlined in this paper offers a means for using a model of the underlying, fundamental behavioral characteristics of the work domain in a principled manner to generate well-grounded decision support concepts for the cognitive demands facing the human-machine decision-making team. This type of approach is essential to bridge the gap between a cognitive analysis of the work domain and the development of innovative decision aids for "envisioned world" types of problems to provide highly effective and robust decision-making performance.

5.5 ACKNOWLEDGMENT

The authors would like to recognize the support of DARPA's Command Post of the Future project for funding the design and development of the Joint Operational Environment under a contract to Logica Carnegie Group (D88D–17–99–C–0047). Specifically, we would like to thank Ward Page and Tom Cook for their vision of the utility of cognitive systems engineering to the development of command and control systems. We would also like to recognize Bridagier General Keith Holcomb, Colonel Larry Shattuck, and Major Jeff From for their support and military expertise in the development of JOE. Finally, we would like to acknowledge the intellectual contribution of Emily Patterson and David Woods with whom we have had many stimulating discussions. We also wish to thank Gary Klein and Kim Vicente for their thoughtful review and constructive comments of an earlier draft of this chapter.

REFERENCES

Bennett, K. B., & Flach, J. M (1992). Graphical displays. Implications for divided attention, focused attention, and problem solving. *Human Factors*, *34*, 513–533.

Carroll, J. M., Kellogg, W. A., & Rosson, M. B. (1991). The task-artifact cycle. In J. M. Carroll (Ed.), *Designing interaction: Psychology at the human-computer interface* (pp. 74–102). Cambridge, England: Cambridge University Press.

Chalmers, B. A., Easter, J. R., & Potter, S. S. (2000, June). Decision-centred visualisations for tactical decision support on a modern frigate. In *Proceedings of 2000 Command and Control Research and Technology Symposium, Making Information Superiority Happen*. Monterey, CA: Naval Postgraduate School.

Cook, R. I., Woods, D. D., Walters, M., & Christoffersen, K. (1996, August). Coping with the complexity of aeromedical evacuation planning: Implications for the development of decision support systems. In *Proceedings of the 3rd Annual Symposium on Human Interaction with Complex Systems*, Dayton, OH: IEEE.

Dekker, S., & Woods, D. D. (1997, October). *The envisioned world problem in cognitive task analysis*. Paper presented at the ONR/NATO Workshop on Cognitive Task Analysis, Washington, DC

Gualtieri, J. W., Roth, E. M., & Eggleston, R. G. (2000). Utilizing the abstraction hierarchy for role allocation and team structure design. In *Proceedings of HICS 2000—5th International Conference on Human Interaction with Complex Systems* (pp. 219–223). Urbana-Champaign, IL: IEEE.

Lind, M. (1991). Representations and abstractions for interface design using multilevel flow modelling. In G. R. S. Weir & J. L. Alty (Eds.), *Human-computer interaction and complex systems.* London: Academic Press.

Lind, M. (1993, July). Multilevel flow modeling. *AAAI93 Workshop on Reasoning about Function*, Washington, DC.

Logica Carnegie Group: (2000a). *Command Post of the Future: Cognitive modeling and engineering* (1999 Year End Report). Aberdeen Proving Ground, MD: U.S. Army Research Laboratory, Human Research and Engineering Directorate.

Potter, S. S., Ball, R. W., Jr., & Elm, W. C. (1996, August). Supporting aeromedical evacuation planning through information visualization. In *Proceedings of the 3rd Annual Symposium on Human Interaction with Complex Systems* (pp. 208–215). Dayton, OH: IEEE.

Potter, S. S., McKee, J. E., & Elm, W. C. (1997). *Decision centered visualization for the military capability spectrum project* (Unpublished technical report). Pittsburgh, PA: Carnegie Group.

Potter, S. S., Roth, E. M., Woods, D. D. & Elm, W. (2000). Bootstrapping multiple converging cognitive task analysis techniques for system design. In J. M. C. Schraagen, S. F. Chipman, & V. L. Shalin (Eds.), *Cognitive task analysis*. Mahwah, NJ: Erlbaum.

Potter, S. S., Woods, D. D., Hill, T., Boyer, R. L., & Morris, W. S. (1992, October). Visualization of dynamic processes: Function-based displays for human-intelligent system interaction. In *Proceedings of the 1992 IEEE International Conference on Systems, Man, and Cybernetics* (pp. 1504–1509). Chicago, IL: IEEE.

Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to cognitive engineering.* New York: Elsevier.

Rasmussen J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering.* New York: Wiley.

Reising, D. V., & Sanderson, P. (1998). Designing displays under ecological interface design: Towards operationalizing semantic mapping. In *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting* (pp. 372–376). Santa Monica, CA: HFES.

Roth, E. M., Lin, L., Thomas, V. M., Kerch, S., Kenney, S. J., & Sugibayashi, Nubuo (1998). Supporting situation awareness of individuals and teams using group view displays. In *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting* (pp. 244–248). Santa Monica, CA: HFES.

Roth, E. M., Lin, L., Kerch, S., Kenney, S. J., & Sugibayashi, N. (In press). Designing a first-of-a kind group view display for team decision making: A case study. In E. Salas & G. Klein (Eds), *Linking expertise and naturalistic decision making.* Mahwah, NJ: Erlbaum.

Roth, E. M., Malsch, N., Multer, J., & Coplen, M. (1999). Understanding how train dispatchers manage and control trains: A cognitive task analysis of a distributed planning task. In *Proceedings of the Human Factors and Ergonomics Society 43rd Annual Meeting* (pp. 218–222). Santa Monica, CA: HFES.

Roth, E. M., Malin, J. T. , & Schreckenghost, D. L. (1997). Paradigms for intelligent interface design. In M. Helander, T. Landauer, & P. Prabhu (Eds.), *Handbook of human-computer interaction* (2nd ed., pp. 1177–1201). Amsterdam: North-Holland.

Roth, E. M., & Mumaw, R. J. (1995). Using cognitive task analysis to define human interface requirements for first-of-a-kind systems. In *Proceedings of the Human Factors and Ergonomics Society 39th Annual Meeting* (pp. 520–524). Santa Monica, CA: HFES.

Rouse, W., & Morris, N. (1986). On looking into the black box: Prospects and limits in the search for mental models. *Psychological Bulletin*, *100* (3), 349–363.

Schraagen, J.M.C., Chipman, S.F., & Shalin, V.L. (Eds.) (2000). *Cognitive task analysis*. Mahwah, NJ: Erlbaum.

Simon, H. (1981). *The sciences of the artificial*. Cambridge, MA: MIT Press.

Smith, P., Woods, D. D., McCoy, E., Billings, C., Sarter, N., Denning, R., & Dekker, S. (In press). Using forecasts of future incidents to evaluate future ATM system designs. *Air Traffic Control Quarterly.*

Vicente, K. (1999). *Cognitive work analysis.* Mahwah, NJ: Erlbaum.

Vicente, K. J., Christoffersen, K., & Hunter, C. N. (1996). Response to Maddox critique. *Human Factors, 38* (3), 546–549.

Vicente, K. J., & Rasmussen, J. (1990). The ecology of human-machine systems II: Mediating direct perception in complex work domains. *Ecological Psychology, 2,* 207–249.

Vicente, K., & Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. *IEEE Transactions on Systems Man and Cybernetics, 22*, 589–606.

Woods, D. D. (1988). *The significance messages concept for intelligent adaptive data display* (Unpublished technical report). Columbus, OH: The Ohio State University.

Woods, D. D. (1995). Toward a theoretical base for representation design in the computer medium: Ecological perception and aiding human cognition. In J. Flach, P. Hancock, J. Caird, & K. Vicente (Eds.), *Global perspectives on the ecology of human-machine systems* (pp. 157–188). Hillsdale, NJ: Erlbaum.

Woods, D. D. (1998). Designs are hypotheses about how artifacts shape cognition and collaboration. *Ergonomics*. *41*, 168–173.

Woods, D. D., & Hollnagel, E. (1987). Mapping cognitive demands in complex problem-solving worlds. *International Journal of Man-Machine Studies, 26*, 257–275.

Woods, D. D., & Roth, E. M. (1988). Cognitive engineering: Human problem solving with tools. *Human Factors, 30* (4), 415–430.